

Traverse Planning Experiments for Future Planetary Surface Exploration

S. J. Hoffman, Ph.D¹, S. A. Voels, Ph.D¹, R. P. Mueller² and P. C. Lee, Ph.D³

¹ Science Applications International Corporation, 2450 NASA Parkway, Houston, Texas 77058; PH 281.483.9264 or 281.244.8171; email: stephen.j.hoffman@saic.com or stephen.a.voels@saic.com

² NASA Kennedy Space Center, M6-0399:3415, Kennedy Space Center, Florida 32899; PH 321.867.2557; email: rob.mueller@nasa.gov

³ Mars Institute., NASA Ames Research Center, MS 245-3, Moffett Field, CA 94035; PH 800.792.9813, Extension 707; email: pascal.lee@marsinstitute.net

ABSTRACT

This paper describes the results of a recent (July-August 2010 and July 2011) planetary surface traverse planning experiment. The purpose of this experiment was to gather data relevant to robotically repositioning surface assets used for planetary surface exploration. This is a scenario currently being considered for future human exploration missions to the Moon and Mars. The specific scenario selected was a robotic traverse on the lunar surface from an outpost at Shackleton Crater to the Malapert Massif. As these are exploration scenarios, the route will not have been previously traversed and the only pre-traverse data sets available will be remote (orbital) observations.

Devon Island was selected as an analog location where a traverse route of significant length could be planned and then traveled. During the first half of 2010, a team of engineers and scientists who had never been to Devon Island used remote sensing data comparable to that which is likely to be available for the Malapert region (eg., 2-meter/pixel imagery, 10-meter interval topographic maps and associated digital elevation models, etc.) to plan a 17-kilometer (km) traverse. Surface-level imagery data was then gathered on-site that was provided to the planning team. This team then assessed whether the route was actually traversable or not.

Lessons learned during the 2010 experiment were then used in a second experiment in 2011 for which a much longer traverse (85 km) was planned and additional surface-level imagery different from that gathered in 2010 was obtained for a comparative analysis.

This paper will describe the route planning techniques used, the data sets available to the route planners and the lessons learned from the two traverses planned and carried out on Devon Island.

INTRODUCTION

This paper describes the results from several field tests of an approach for planning and carrying out a robotic traverse across many tens to hundreds of kilometers of a planetary surface. These tests included not only the processes to be

used, but considered the minimum data sets necessary to support these processes. The traverse scenario was focused somewhat narrowly on relocating assets (e.g., rovers, habitats, power sources, etc.) necessary for human crews to explore lunar sites for relatively short periods of time. The experiment team hypothesized that remotely sensed data – such as panchromatic and monochromatic imagery, elevation and radar – with a resolution no better than that which will be available from the Lunar Reconnaissance Orbiter (LRO) will be sufficient for planning these traverses and then carrying out the teleoperated movement of these large and valuable assets.

After describing some background to help explain the origin of questions that lead to this experiment, the paper will discuss details associated with a 2010 field test and a follow-on test in 2011. Both tests were carried out near the Haughton Mars Project Research Station (HMPRS) located on Devon Island in northern Canada, a site with terrain features comparable to those expected to be found at sites of interest on the Moon and for which appropriate data sets are available.

Background. In late 2008 it became clear that many space agencies were engaged in plans and preparations for missions beyond low-Earth orbit that could benefit from early coordination. As a result, the International Space Exploration Coordination Group (ISECG) was formed. In early 2009 the ISECG endorsed development of a Reference Architecture for Human Lunar Exploration and invited interested agencies to participate (Anon., 2009). To further the goal of cooperation, the ISECG established the International Architecture Working Group (IAWG) and the International Objectives Working Group (IOWG) to analyze the lunar exploration objectives of participating agencies. A series of workshops conducted during 2009 led to the development of a set of common lunar exploration goals. These goals were accepted by the ISECG in December 2009. A set of strategic guidelines was combined with these goals to guide another working group in the development of a reference architecture that was adopted by the ISECG (Anon., July 2010).

This reference architecture is composed of four phases during which a range of international human-rated and robotic technologies will be deployed over time on the lunar surface. It also provides for continuous robotic and human exploration activity in multiple locations on the Moon. The four phases include:

(1) *Robotic precursor phase*: Provides early technology demonstrations and engagement among international partners, the scientific community, and the public. It highlights important activities intended to reduce the risks associated with human missions and to ensure sustainability of the architecture.

(2) *Polar exploration and system validation phase*: Initiates human exploration of the Moon. It leverages the robotic precursor work to deploy and test an international fleet of crewed rovers and supporting robots in preparation for more aggressive human and robotic lunar exploration.

(3) *Polar relocation phase*: Relocates the fleet of robots and rovers, controlled from Earth, from the pole to new sites of interest. Along the way, this fleet will perform scientific studies and enable interactive participation from the public. Once in place, the robots and rovers meet and assist human crews landing at each new site.

(4) *Nonpolar relocation and long-duration phase*: May involve multiple short missions to various lunar sites of interest or long-duration missions of about 70 days at one site.

A specific scenario developed during formulation of the “polar relocation phase” of this reference architecture – an approximately 200-kilometer traverse from the rim of Shackleton Crater to the Malapert Massif – was used to guide the development of the field tests described in this paper. These field tests did not replicate the Shackleton-to-Malapert traverse in detail but took into account those features of the traverse (e.g., transitioning across rolling terrain, using rock outcrops as waypoints and scientific targets, transitioning around the rim of a crater, etc.) that were noted by the planners as they considered alternative means of completing their representative “polar relocation.”

Three specific issues related to the traverse planning and execution aspect of this scenario were selected for analog field tests (Note: NASA/TM-2011-216157 contains a complete discussion of these tests and the results):

- (1) *The available remote observation data sets of a region to be traversed are sufficient for planning primary and alternate traverse routes* (named TRPF-H1).
- (2) *Robotically implemented traverse route execution will require surface-level imagery to identify and maneuver around local hazards/obstacles* (named TRPF-H2).
- (3) *Route-traversing efficiency will improve in direct proportion to the number of surface-level imagery sources used to support a traverse* (named TRPF-H3).

EXPERIMENT IMPLEMENTATION

A test area with the following characteristics was needed to provide a reasonable simulation environment in which to conduct the traverse tests: (a) relatively large area (many square kilometers), preferably free of vegetation, human-built obstructions (e.g., fences, roads, etc.), or other ground cover (e.g., snow, ice, etc.); (b) terrain features and topography representative of the lunar South Pole region being simulated in these tests; and (c) areas in which a traverse could be successfully carried out as well as areas in which these activities could clearly not be carried out to determine the test teams’ ability to distinguish between the two. The region around the HMPRS on Devon Island, Nunavut, Canada, possessed these features and therefore was selected for these tests. The remaining activities in this experiment were implemented in four parts:

Part 1 *Definition of the analog traverse waypoints by the principal investigator (PI), the available remote observation datasets, and other constraints for planning the traverse.* The PI established a start point and an endpoint for the simulated traverse; these two points are separated by a straight-line distance of just over 17 kilometers (11 miles; ~15 percent of the straight-line distance between Shackleton Crater and Malapert Massif). Eight waypoints were inserted between the start point and endpoint of the simulated traverse. These waypoints are meant to simulate points of interest that could have been introduced for science reasons or operations reasons (e.g., a high point in which solar illumination was available). But the process by which these waypoints would have been selected was not part of the simulation and therefore will not be further discussed. (Note: There is also a practical reason for including these waypoints: to provide a prominent landmark for the person on the ground to find independent of latitude/longitude should there be any issue with

map/image ties with the remote observation data sets being used.) The cumulative straight-line distance between all eight of these waypoints is 22.7 kilometer.

The following datasets were selected by the PI for use by the planning team: (a) satellite and air photographs (equivalent to ~1 to 2 meter/pixel); (b) topographic maps (10-meter elevation contour) and digital elevation model (DEM) maps (selected areas); and (c) satellite radar imaging (25 meter/pixel resolution). All of these data were imported into Google Inc.'s Google® Earth (GE) tool as the general planning environment.

Part 2 *Selection of a traverse route and required ground observations by a route planning team (six people).* The traverse path and photographs to be taken were done in a four-step process.

Step 1: Using material (i.e., remote observation data) provided to the planning team, team members were asked to determine the shortest route that could be traversed by all elements of the convoy. If the planning team was fairly sure, but not completely convinced, that the particular route was feasible, team members were asked to prepare a possible alternate route to bypass the questionable area. The planning team documented the routing as a set of additional waypoints between those selected by the PI. The group was also asked to assign a trafficability rating to the selected route using a set of defined categories.

Step 2: With the traverse path determined, the planning team was asked to decide which ground-level photographs along the primary route are needed to confirm that the path picked is passable. The planning team was advised to take into account turns, possibly hidden depressions, and rocky terrain in determining the ground-level photographs needed. The group documented these photograph requests as a set of locations (latitude, longitude), bearing, and field of view (FOV) size for each photograph.

Step 3: The team was asked to revisit the path selected and determine whether there are locations in which use of a point of perspective would be useful (e.g., to check an alternate route or to provide a view not from the traverse path). Again, the team was asked to document these photograph requests as a set of locations (latitude, longitude), bearing, and FOV size for each photograph.

Step 4: Iterate through the first three steps until a consensus was reached. This was not as much a planned step as an immediate lesson learned, in that the planning team members put forward several different paths, each considered "reasonable" or "safe" by one or more of the team members. Discussion of the rationale behind these alternatives led to an iterative process to reach a consensus path. Figure 1 shows the eight-waypoint traverse route, with some of the local deviations visible (better detail is available in Hoffman, et.al., 2011).

Part 3 *Gathering of the ground observations on Devon Island.* The third part of the experiment consisted of the PI traveling the route laid out by the planning team to gather the requested photographs and document the ground truth. The PI obtained the photographs, tagged them with Global Positioning System (GPS) coordinates, and transmitted them from the HMPRS base camp to NASA's Future Flight Central (FFC) facility at the Ames Research Center (ARC), which was used by the planning team during Part 4 of the experiment. The PI also noted and documented ground-truth data regarding terrain characteristics along the planned route, in particular whether



Figure 1. Traverse route for the 2010 field experiments

the planned route was “safe” for all vehicles to follow considering rocks or slope as well as the trafficability of the route.

When the PI believed the route was unsafe for the convoy or thought the traverse team may perceive a problem, the PI took additional photographs that might help the groups resolve the situation. These

additional photographs were only taken as a time-saving measure, anticipating the

team would request such a photograph during the execution phase. During the execution phase, the traverse team would only be allowed to view the photographs if the team members requested a photograph that matched one of the photographs the PI took.

Part 4 *Simulation of the planned traverse using the data gathered on Devon Island (but previously transmitted and stored at ARC due to communications limitations) and any real-time data requests.* The final part of the experiment was to “execute” the traverse. The images and GE view along the planned route on Devon Island were presented to the traverse teams on displays in the FFC (see Figure 2); the primary image was also shown on WebEx® (Cisco Technology, Inc.) for remote team member(s). The images were displayed no faster than the nominal ground-ruled speed of the traverse (i.e., 5 km/h, which is 72 seconds per image for images 100 meters apart). For each image displayed, the traverse team decided and recorded whether the route was passable for the segment represented by the image and determined the terrain trafficability category. Because slope is part of trafficability, the traverse team was also provided with the ground-truth slope at the position of the photograph.

At locations at which the planning team had previously determined an alternate route might be needed, the traverse team was presented images of both the primary and the alternate route from that location. The traverse team then determined whether the two images provided sufficient information to select one route over the other. In cases in which there was no clear advantage of one route over the other route, the traverse team was instructed to continue on the primary route.



Figure 2. Simulated remote driving station at NASA ARC FFC. The two images on the right are examples of surface-level imagery sent to the traverse planning team. The image on the left is the GE view of the photograph location with the traverse route overlaid on the surface image.

When the team came across a section of route that they determined was impassable, two scenarios could come into play depending on the situation:

- If team members had already determined that this section might be a problem, the team would have already planned an alternate route. In this case, the team would back up to the start of the alternate route and proceed on that alternate route.
- If there was no alternate route or all alternate routes had been exhausted, the team members were asked to plan an alternate route “on the fly.” Originally, the plan had been to allow the traverse team to request additional imagery at the “current” traverse location. The PI on Devon Island would have taken this additional imagery in real time and sent it back in near-real time. However, weather conditions at Devon Island did not allow for this scenario. Consequently, the test director in the FFC used the ground-truth notes and additional imagery previously sent by the PI to make rulings on proposed alternate routes.

The time spent evaluating and discussing each imagery location and the discussions concerning alternate routes were recorded by the test director.

RESULTS AND CONCLUSIONS FOR THE 2010 FIELD TESTS

During this experiment, and the analysis and discussion of the data from it, the team made the following observations and came to the following conclusions:

TRPF-H1. Remote observation data are insufficient to plan a definitive route that is completely free of undetected obstacles that would stop vehicles of the type assumed here. Sections of the planned route were impassable. The primary type of obstacle encountered was steep terrain – cliffs, large outcrops, steep slopes, etc. (see photographs on the right of Figure 2). While remote observation data were marginal for detection of obstacles such as rocks, there were no sections that were impassable due to rocks alone.

TRPF-H2. Surface-level imagery, which provided a significant improvement in the knowledge of the route, was required to complete the traverse. All of the major obstacles were identifiable in the surface-level imagery. While some of marginal cases were not immediately seen, they likely would have been detected had real-time, full-motion, surface-level imagery been available. As with the remote observation data, determination of slopes was the most difficult problem to detect and evaluate.

TRPF-H3. The original proposed field test involving multiple views was not implemented due to several issues in the field (see TRPF-H3 introduction). However, based on some limited opportunistic data, it appears there is a concern that it may take too long for the remote “drivers” to correlate multiple views. Additional work needs to be done on exactly how multiple views would be used and how imagery can be quickly correlated. But given the overall results from this opportunistic experiment, any additional work in this area should be considered low priority.

General Observations. Although the surface imagery was sufficient to carry out the experiment, the “driving” and evaluation of the route would have been much easier if additional information had been directly available to help interpret the photographs. Examples include the following:

(a) Distance or size cues are clearly needed in the photographs. This could be achieved either by projecting additional information on the photographs or having or projecting known cues into the scene when the imagery is taken. Use of video instead of still imagery (an implementation limitation in this experiment) would have given some additional distance/size cues, but it is unclear whether those cues would have been sufficient. Another approach would be to use a scanning range-finding device (e.g., a laser) to determine distance and then overlay the results on the imagery.

(b) Slope measurements need to account for cross-track slopes. The original planning analysis only calculated slopes along the direction of travel. The ground truth showed that there were several locations in which the cross-track slope was significant – in some cases, much steeper than the slope along the direction of travel.

(c) There needs to be a more automated mechanism to display imagery and additional metadata from the remote site that the team executing the traverse is allowed to see. For example, the slopes measured at the surface-level imagery location should be displayed with the photographs. Similarly, better software support for the use of multiple screens would allow both the current and the previous photograph to be displayed and changed with a single command. Overlay of the route on the photograph would be useful. The vast majority of the routes should have been straight ahead for the next segment. The evaluators mentioned that the overlay would still be useful to see upcoming turns past the next point.

(d) Acknowledging that the use of multiple still images for this experiment was the result of logistical constraints, if this approach were to be used for an actual traverse, the traverse team determined that a distance of 100 meters between surface imagery locations was too large. The team was usually looking or concentrating on the next 25 to 50 meters. While the photographs were taken at a resolution that allowed zooming, the additional time needed to zoom in the images was usually not taken. Several locations on hills (especially the descending side) and even flat areas with minor local rises (2 to 3 meters) led to part of the route not being visible in the photograph. Although ground truth showed these were not problem areas, it did lead

to loss of situational awareness and increased evaluation times. Thus, given the bandwidth limitation for this experiment and the manner in which the teams used the photographs, more frequent but lower-resolution images would have probably worked better than less frequent but higher-resolution images.

FOLLOW-UP FIELD TEST IN 2011

During 2011 an opportunity arose to apply some of the lessons learned from the 2010 traverse experiment as well as gather additional data to help answer several of the identified difficulties. An AM General LLC Humvee® (or formally high-mobility multipurpose wheeled vehicle [HMMWV]) was to be moved from the west coast of Devon Island to the HMPRS, a distance of approximately 75 kilometers. A route was needed for this traverse across this area that had not been previously visited by anyone on the driving team. Three objectives were set for this 2011 experiment, each associated with a lesson learned or an observation from the 2010 experiment:

- (1) Continue to refine the traverse planning process by including additional data (i.e., the slope map) and by planning longer routes through unexplored terrain.
- (2) Obtain ground-truth imagery at intervals much less than the 100-meter interval used during 2010 to understand the improvement in (a) situational awareness during a supervised robotic traverse and (b) to improve local replanning should an obstacle or hazard be encountered.
- (3) Understand the impacts of obtaining surface-level imagery from a location higher than and lower than that used in the 2010 experiment.

Data sets used for planning this traverse were comparable to those used in 2010 but with the addition of a slope map. This slope map used the local DEM data to calculate the maximum slope at each point in this region. The slope data became another factor in selecting a “safe” route. The resulting traverse route is shown in Figure 3. This route had a total distance of approximately 90 kilometers. One potential problem was noted near the beginning of the traverse, resulting in a planned alternate route in that area.

Four cameras were mounted on the Humvee (see Figure 4) – one camera at approximately eye level to produce photographs from a perspective comparable to that of the 2010 experiment, one camera somewhat lower level to approximate a small robotic scout, and a pair of cameras on top of the Humvee for a perspective approximating the view from the top of a small pressurized rover as well as providing a three-dimensional view. Photographs were taken at five-second intervals so that at a nominal speed of 10 - 15 km/h the photograph data set would be at an interval of approximately 15 – 20 meters.

An all-terrain vehicle (ATV) performed the role of a small scout rover to (a) locate a precise route between waypoints, (b) help locate soft ground, and (c) provide the suggested size cue in photographs. This ATV was typically 100 meters in front of the Humvee. A suspected obstacle early in the traverse was in fact real, so the alternate path was used. In addition, five additional unidentified obstacles were encountered during the remainder of the traverse. All of these obstacles were slope features that would have been too steep for the criteria stated in the original experiment. A safe path around these features was typically located within a few tens

of meters from the baseline path. A significant amount of soft ground caused by snowmelt was also encountered but was not recorded as an obstacle. The total distance driven, including all deviations from slope and soft-ground obstacles, was 102 kilometers, a 16 percent increase over the planned route.

A total of 65,544 photographs were collected from all four cameras, with an

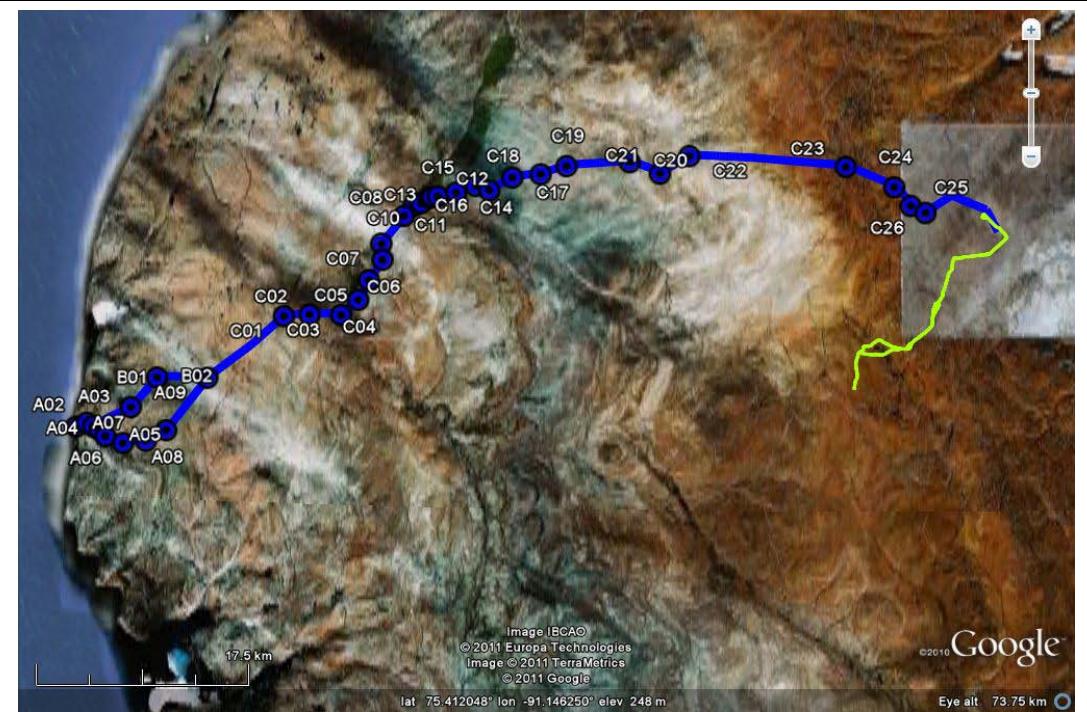


Figure 3. The 2011 traverse (85 kilometer; shown in blue) with the 2010 traverse (primary and alternates shown in yellow and green, respectively) for comparison.

additional 3,232 context photographs from other cameras. Of these 24,361 photographs were collected while the vehicles were moving and form the database from which comparative analyses identified in the 2010 experiment will be made. These analyses are still under way. However, several conclusions can be drawn from the results already gathered.

For the TRPF-H1 hypothesis, the data set used was sufficient to plan a traverse without encountering an obstacle that could not be maneuvered around within a few tens of meters. The number of unidentified obstacles encountered was somewhat lower per kilometer of planned traverse but still not zero. The revised version of the slope map was somewhat useful but the resolution of the data from which it was constructed needs to be significantly higher (on par with the image resolution) for this to be a useful planning tool.

For the TRPF-H2 hypothesis, the increased frequency of the photographs definitely improved situational awareness and a five-second interval may be sufficient for this type of teleoperated driving when compared to continuous video. Having an object of known size, such as the ATV, in the field of view also provides a much better sense of scale, slope and distance. The advantages and disadvantages of high versus low camera positions and stereo versus mono images are still under analysis so no conclusions can be offered.

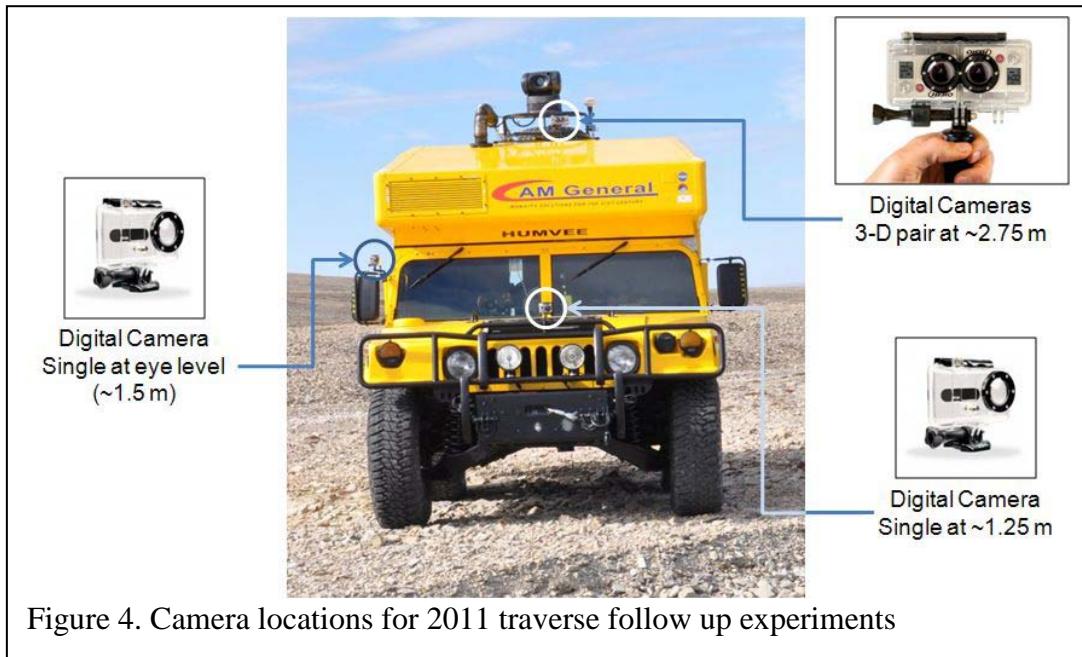


Figure 4. Camera locations for 2011 traverse follow up experiments

CONCLUSION

Several analog experiments have been conducted to better understand the process and required data sets needed for a teleoperated robotic traverse across planetary surfaces as proposed by the ISECG. Results to date indicate that remotely sensed data sets comparable to those already in hand for the Moon (and to a more limited extent for Mars) coupled with surface-level imagery should also be sufficient for actually driving planned traverse routes. Characteristics of the surface-level imagery required is still being assessed, but preliminary results indicate that discrete images obtained at a frequency to update the situational view every 10 to 20 meters should be sufficient.

REFERENCES

- Anon. (March 10-12, 2009). "Advancing the Global Exploration Strategy Human Exploration of the Moon." Summary of scenario discussions held by: International Space Exploration Coordination Group; Yokohama, Japan.
- Anon. (July 2010). The ISECG Reference Architecture for Human Lunar Exploration – Summary Report. Available at:
<http://www.globalspaceexploration.org/home;jsessionid=00B68432DE6CEF02E137B2559276C6B1>. Accessed September 16, 2011.
- Hoffman, Stephen J., Leonard, Matthew J., and Lee, Pascal (September 2011). "Evaluation of Robotic Systems to Carry Out Traverse Execution, Opportunistic Science, and Landing Site Evaluation Tasks," NASA/TM-2011-216157, NASA Johnson Space Center, Houston, Texas.